

Lempel-Ziv Computation In Compressed Space (LZ-CICS[®])

Dominik Köppl¹ and Kunihiko Sadakane²

¹Department of Computer Science, TU Dortmund, Germany

²Graduate School of Information Science and Technology, University of Tokyo, Japan

Abstract

We show that both the Lempel-Ziv-77 and the Lempel-Ziv-78 factorization of a text of length n on an integer alphabet of size σ can be computed in $\mathcal{O}(n \lg \lg \sigma)$ time (linear time if we allow randomization) using $\mathcal{O}(n \lg \sigma)$ bits of working space. Given that a compressed representation of the suffix tree is loaded into RAM, we can compute both factorizations in $\mathcal{O}(n)$ time using $z \lg n + \mathcal{O}(n)$ bits of space, where z is the number of factors.

1 Introduction

The Lempel-Ziv-77 (LZ77) [1] and the Lempel-Ziv-78 (LZ78) [2] factorization divide a text into factors that capture repetitions in the text. Although both factorizations are found in major text processing tools like compressors or full text indices, computing any of the two factorizations is a bottleneck in terms of space and time. In practice, compressors still use a sliding window or discarding techniques to avoid high resource consumption. Hence, one might ask whether it is possible to lower the space bound in the light of recent approaches in the field of succinct data structures, while still allowing linear running time.

In this article, we show that the LZ77 and the LZ78 factorization of a text of length n on an integer alphabet of size σ can be computed

- with $\mathcal{O}(n \lg \sigma)$ bits of working space in either $\mathcal{O}(n)$ randomized or $\mathcal{O}(n \lg \lg \sigma)$ deterministic time, and
- with $z \lg n + \mathcal{O}(n)$ additional bits of working space in linear time, given that we have access to the compressed suffix tree of the text.

2 Related Work

We are aware of the following results for LZ77: The currently most space efficient algorithm is due to Kosolobov [3] whose algorithm runs in $\mathcal{O}(n(\lg \sigma + \lg \lg n))$ time and uses only εn bits of working space, provided that we have read-access to the text. A trade-off algorithm is given by Kärkkäinen et al. [4], using $\mathcal{O}(n/d)$ words of working space and $\mathcal{O}(dn)$ time. By setting $d \leftarrow \log_\sigma n$ we get $\mathcal{O}(n \lg \sigma)$ bits of working space and $\mathcal{O}(n \log_\sigma n)$ time. The algorithm of Belazzougui and Puglisi [5] derives its dominant terms in space and time from the same data structure [6] as we do; the LZ77 factorization algorithms of both papers work with the same space and time bounds. Their algorithm uses only the Burrows-Wheeler transform (BWT) [7] construction algorithm from [6]. By exchanging it with an improved version [8], we expect that their algorithm will run in deterministic linear time.

Since LZ78 factors are naturally represented in a trie, the so-called **LZ trie**, improving LZ78 computation can be done, among others, by using sophisticated trie implementations [9, 10], or by superimposing the suffix tree with the suffix trie [11, 12]. We follow the latter approach. There, both Nakashima et al. [11] and Fischer et al. [12] presented a linear time algorithm, using $\mathcal{O}(n \lg n)$ and $(1 + \varepsilon)n \lg n + \mathcal{O}(n)$ bits of space, respectively.

Based on the data structures of the compressed suffix tree, we derive our techniques from an approach [12] using the suffix tree topology with succinct representations of the suffix array, its inverse and the longest common prefix array. For both factorization variants, Fischer et al. [12] store the inverse suffix array and parts of the enhanced suffix array in $(1 + \varepsilon)n \lg n + \mathcal{O}(n)$ bits of space such that they can access leaves of the suffix tree in text order, and can compute the string depth of internal nodes, both in constant time. Unlike the here presented approach, their algorithms overwrite the working space multiple times, and use a complicated counting for the LZ78 trie nodes.

3 Preliminaries

Our computational model is the word RAM model with word size $\Omega(\lg n)$ for some natural number n . Accessing a word costs $\mathcal{O}(1)$ time. We assume that the function $\text{popcount}(w)$, counting the set bits in a word w , can be computed in constant time. Otherwise, we build a lookup-table [13] supporting popcount with two rank queries in constant time. The lookup-table fits into our working space.

Let Σ denote an integer alphabet of size $\sigma = |\Sigma| \leq n$. We call an element $T \in \Sigma^*$ a **string** or **text**. Its length is denoted by $|T|$. The empty string is ϵ with $|\epsilon| = 0$. We access the j -th character of T with $T[j]$ for $1 \leq j \leq |T|$. Given $x, y, z \in \Sigma^*$ with $T = xyz$, we call x , y , and z a **prefix**, a **substring**, and a **suffix** of T , respectively. In particular, the suffix starting at position j of T is called the **j -th suffix** of T .

We call strings on the binary alphabet $\{0, 1\}$ **bit vectors**. For a bit vector B , we are interested in answering the following queries for $c \in \{0, 1\} \cup \{0, 1\}^2$:

- $B.\text{rank}_c(j)$ counts the number of ‘ c ’s in $B[1, j]$, and
- $B.\text{select}_c(j)$ gives the position of the j -th ‘ c ’ in B .

We can answer both types of queries due to a result of Raman et al. [14]: There is a data structure taking $o(|B|)$ extra bits of space to answer rank and select queries in constant time. It can be constructed in time linear to $|B|$. We say that a bit vector has a **rank-support** and a **select-support** if it provides constant time access to rank and select, respectively.

The zero-order entropy $H_0(n, z)$ of a bit vector of length n storing z ones is defined by $H_0(n, z) = z \lg(n/z) + (n - z) \lg(n/(n - z))$. Such a bit vector can be compressed such that it consumes $H_0(n, z) + o(n)$ bits, supporting access and rank in constant time (e.g., [14]).

In the rest of this paper, we take a read-only text T of length n , which is subject to the LZ77 or the LZ78 factorization. Let $T[n]$ be a special character appearing

nowhere else in T , so that no suffix of T is a prefix of another suffix of T . Without loss of generality, we assume that Σ is the *effective* alphabet of T , i.e., each character of Σ appears in T at least once. Otherwise, we can reduce the alphabet to the effective alphabet by sorting the characters with a linear time integer sorting algorithm using $n \lg \sigma + \mathcal{O}(1)$ working space [15], and an array with $\sigma \lg \sigma$ bits to reconstruct the former alphabet.

3.1 Lempel-Ziv Factorization

A **factorization** of T with size z partitions T into z substrings $T = f_1 \cdots f_z$. These substrings are called **factors**. In particular, we have:

Definition 3.1. A factorization $f_1 \cdots f_z = T$ is called the **LZ77 factorization** of T iff $f_x = \operatorname{argmax}_{S \in S_j(T) \cup \Sigma} |S|$ for all $1 \leq x \leq z$ with $j = |f_1 \cdots f_{x-1}| + 1$, where $S_j(T)$ denotes the set of substrings of T that start strictly before j (for $1 \leq j \leq |T|$).

Different to the LZ77 factorization, the classic-LZ77 factorization adds an additional character to the referencing factors:

Definition 3.2. A factorization $f_1 \cdots f_z = T$ is called the **classic-LZ77 factorization** of T iff f_x is the shortest prefix of $f_x \cdots f_z$ that occurs exactly once in $f_1 \cdots f_x$.

Definition 3.3. A factorization $f_1 \cdots f_z = T$ is called the **LZ78 factorization** of T iff $f_x = f'_x \cdot c$ with $f'_x = \operatorname{argmax}_{S \in \{f_y : y < x\} \cup \{\epsilon\}} |S|$ and $c \in \Sigma$ for all $1 \leq x \leq z$.

3.2 Suffix Tree

The **suffix trie** of T is the trie of all suffixes of T . The **suffix tree (ST)** of T , denoted by ST , is the tree obtained by compacting the suffix trie of T . We denote the root node of ST by root . Our approach uses a compressed representation of ST , consisting of

- the ψ -array [16] with $\text{SA}[i] = \text{SA}[\psi(i)] - 1$ for $1 \leq i \leq n$ with $\text{SA}[i] \neq n$ (and $\psi(i) = \text{SA}^{-1}[1]$ for $\text{SA}[i] = n$), and
- a $4n + o(n)$ -bit balanced parenthesis representation (BP) [17] of the tree topology [18], equipped with the minmax tree [19] for navigation.

By employing the algorithm of Belazzougui [6] on the text, we can build the compressed suffix tree consuming $\mathcal{O}(n \lg \sigma)$ bits of space in either $\mathcal{O}(n)$ randomized time or $\mathcal{O}(n \lg \lg \sigma)$ deterministic time.

Due to the BP representation, each node of the suffix tree is uniquely identified by its pre-order number. A rank- and a select-support on the BP representation enable us to address a node by its pre-order number in constant time. If the context is clear, we implicitly convert an ST node to its pre-order number, and vice versa.

Each leaf is labeled *conceptually* by the text position where its corresponding suffix starts (see Figure 1). We write $\text{label}(\ell)$ for the label of a leaf ℓ . Reading the leaf labels in depth first order returns the suffix array, which we denote by SA . We do neither store SA nor the leaf labels.

Algorithm 1 Implementation of *str_depth*

Require: suffix tree node v

```
1: if  $v$  is an internal node then
2:    $\ell \leftarrow \text{least\_leaf}(v.\text{child}(1))$ 
3:    $\ell' \leftarrow \text{least\_leaf}(v.\text{child}(2))$   $\{v.\text{child}(2)$  exists since  $v$  is an internal node $\}$ 
4:    $m \leftarrow 0$ 
5:   while  $\text{head}(\ell) = \text{head}(\ell')$  do
6:      $\ell \leftarrow \text{next\_leaf}(\ell)$ 
7:      $\ell' \leftarrow \text{next\_leaf}(\ell')$ 
8:     incr  $m$ 
9:   end while
10:  return  $m$ 
11: else if  $v$  is a leaf then  $\{\text{works only if } \text{label}(v) \text{ is available}\}$ 
12:  return  $n + 1 - \text{label}(v)$ 
13: end if
```

- $\text{head}(\ell)$ retrieves the first character of the suffix whose starting position coincides with the label of the leaf ℓ . Since Σ is the effective alphabet of T , each character of Σ occurs in T . So **root** has σ children, each corresponding to a different character. Besides, the order of **root**'s children and of the characters of Σ is the same. Hence, $\text{child_rank}(\text{level_anc}(\ell, \text{depth}(\text{root}) + 1)) = \text{head}(\ell)$ holds, and the left hand side can be computed in constant time.
- smallest_leaf selects the leaf with the label 1. By a linear scan over the ψ -array, we can find the value $\alpha := \mathbf{SA}^{-1}[1]$, so that $\psi^k[\alpha] = \mathbf{SA}^{-1}[k+1]$ for $0 \leq k \leq n-1$. We store α to answer smallest_leaf by $\text{leaf_select}(\alpha)$.
- $\text{next_leaf}(\ell)$ selects the leaf labeled with $\text{label}(\ell) + 1$. We can compute it in constant time, since $\text{next_leaf}(\ell) = \text{leaf_select}(\psi[\text{leaf_rank}(\ell)])$.
- $\text{str_depth}(v)$ returns the string depth of an *internal* node. We use the ψ -array and the head -function to compute $\text{str_depth}(v)$ in time proportional to the string depth (see Algorithm 1). Therefore, we take two different children of v (they exist since v is an internal node), and choose an arbitrary leaf in the subtree of each child. So we have two leaves representing two suffixes whose longest common prefix is the string read from the edge labels on the path from the root to the lowest common ancestor (LCA) of both leaves. Our task is to compute the length of this prefix. To this end, we match the first characters of both suffixes by the head -function. If they match, we use ψ to move to the next pair of suffixes, and apply the head -function again. Informally, applying ψ strips the first character of both suffixes (like taking a suffix link). On a mismatch, we find the first pair of characters that does not belong to the path from the root to v (reading the labels of the edges along this path). We return the number of matched characters as the string depth.

		1	2	3	4	5	6	7
LZ77	Factor	a	aa	b	aba	aaba	aba	\$
	Coding	a	1,2	b	3,3	2,4	3,3	\$

		1	2	3	4	5
classic-LZ77	Factor	a	aab	abaa	abaab	a\$
	Coding	a	1,2,b	3,3,a	5,4,b	1,1,\$

		1	2	3	4	5	6	7	8
LZ78	Factor	a	aa	b	ab	aaa	ba	aba	\$
	Coding	a	1,a	b	1,b	2,a	3,a	4,a	\$

Figure 2: We parse the text `aaababaaabaaba$` by both factorizations. The coding represents a fresh factor by a single character, and a referencing factor by a tuple with two entries. For LZ77, this tuple consists of the referred position and the number of characters to copy. For classic-LZ77, the tuple additionally contains a new character. For LZ78, it consists of the referred index and a new character.

4 Common Settings

We identify factors by text positions, i.e., we call a text position j the **factor position** of f_x ($1 \leq x \leq z$) iff the factor f_x starts at position j . A factor f_x may refer to either (LZ77) a previous text position j (called f_x 's **referred position**), or (LZ78) to a previous factor f_y (called f_x 's **referred factor**—in this case y is also called the **referred index** of f_x). If there is no suitable reference found for a given factor f_x with factor position j , then f_x consists of just the single letter $T[j]$. We call such a factor a **fresh factor**. The other factors are called **referencing factors**. Let z_R denote the number of referencing factors. An example is given in Figure 2.

4.1 Scaffold of the Algorithms

Common to our LZ77- and LZ78-factorization algorithms is the traversal of the compressed suffix tree. In more detail, they share a common framework, which we describe in the following by introducing some new keywords:

Witnesses. Witnesses are *internal* nodes that act as signposts for finding (LZ77) the referred position or (LZ78) the referred index of a factor. The *number of witnesses* z_W is at most the number of *referencing factors* z_R . We will enumerate the witnesses from 1 to z_W by a bit vector B_W on the BP of **ST** with a rank_1 -support. So each witness has, along with its pre-order number, a so-called **witness id** (its B_W -rank).

Passes. Like the LZ77 algorithm in [12], we divide our algorithms in several passes. In a pass, we visit the leaves of **ST** in text position order. This is done by using *smallest_leaf* and then calling *next_leaf* successively. The passes differ in how a

leaf is processed. While processing a leaf ℓ , we want to access $label(\ell)$. We can track the label of the current leaf with a counter variable, since we start at the leaf with the label 1.

Corresponding Leaves. We say that a leaf ℓ *corresponds to* the factor f if $label(\ell)$ is the factor position of f . During a pass, we keep track of whether a visited leaf corresponds to a factor. To this end, for each leaf ℓ corresponding to a factor f , we compute the length of f while processing ℓ . This length tells us the number of leaves after ℓ (in *text order*) that do not correspond to a factor. By noting the next corresponding leaf, we know whether the current leaf is corresponding to a factor — remember that a pass selects leaves successively in *text order*, and *smallest_leaf* is always corresponding to the first factor.

Output Space. Given \mathbf{ST} and ψ , we analyze our algorithms for both factorizations with respect to time and working space. We analyze both factorizations under the assumptions that either the output has to be stored explicitly in RAM, or that the algorithm must stream the output sequentially.

4.2 Loaded Data Structures in RAM

We need the ψ -array taking $\mathcal{O}(n \lg \sigma)$ bits, and \mathbf{ST} 's topology consuming $4n + o(n)$ bits. We spend $o(n)$ additional bits to create a $rank_{\circ}$ -support, a $select_{\circ}$ -support, and a $rank_{\lrcorner}$ -support on the BP sequence of \mathbf{ST} for ranking and selecting leaves and internal nodes separately (nodes and leaves in the BP representation are given by the sequence ‘(’ and ‘()’, respectively). Our algorithms do *not* access the text T .

5 LZ77

Given that the above stated data structures are loaded into RAM, we show that the LZ77 factorization can be computed with $(1 + H_0(n, z))n + z \lg n + o(n)$ additional bits of working space when streaming. It is easy to show that we can store the output in RAM with additionally $z \lg n$ bits.

LZ77 Passes. Common to all passes is the following procedure: For each visited leaf ℓ , we perform a leaf-to-root traversal, i.e., we visit every node on the path from ℓ to **root**. But we visit every node at most once, i.e., we stop the leaf-to-root traversal on visiting an already visited node. Therefore, we create a bit vector B_V with which we mark a visited node. This bit vector is cleared before a pass starts. Since \mathbf{ST} contains at most $n - 1$ internal nodes, a pass can be conducted in linear time.

We perform two passes:

- (a) create B_W in order to determine the witnesses (see Algorithm 2), and
- (b) stream the output by using an array mapping witness ids to text positions (see Algorithm 3).

Pass (a). We follow the approach from [12]. Determining the witnesses is done in the following way: Reaching the root from a leaf corresponding to a factor (while visiting only non-marked nodes) means that we found a fresh factor. Otherwise, assume that we visit an already visited node $u \neq \mathbf{root}$ from a leaf ℓ . If ℓ

Algorithm 2 LZ77 Pass (a)

```
1:  $\ell \leftarrow \text{smallest\_leaf}$ 
2:  $p \leftarrow 1$  {tracks next leaf corresponding to a factor}
3: repeat
4:    $v \leftarrow \text{parent}(\ell)$ 
5:   while  $v \neq \text{root}$  do
6:     if  $B_V[v] = 1$  then {already visited?}
7:       if  $\text{label}(\ell) = p$  then {if the current leaf corresponds to a factor}
8:          $B_W[v] \leftarrow 1$  {then this node is a witness}
9:          $p \leftarrow p + \text{str\_depth}(v)$  {determine next starting factor}
10:      end if
11:      break{on finding a visited node we stop}
12:    end if
13:     $B_V[v] \leftarrow 1$  {visit the node}
14:     $v \leftarrow \text{parent}(v)$  {move upwards}
15:  end while
16:  if  $v = \text{root}$  then { $\ell$  corresponds to a fresh factor}
17:    incr  $p$  {factor is a single character}
18:  end if
19:   $\ell \leftarrow \text{next\_leaf}(\ell)$ 
20: until  $\ell = \text{smallest\_leaf}$ 
```

corresponds to a factor f , u *witnesses* the referred position of f . This means that there is a suffix starting before $\text{label}(\ell)$ having a prefix equal to the string read from the edge labels on the path from the root to u . Moreover, u is the lowest node in the set $\{\text{LCA of } \ell \text{ and } \ell' : \text{label}(\ell') < \text{label}(\ell)\}$ comprising the lowest common ancestors of ℓ with all already visited leaves. So the factor corresponding to ℓ has to refer to a text position coinciding with the label of a leaf belonging to u 's subtree. In order to find the referred position in the next pass, we mark u in B_W . Additionally, we compute the length of f with $\text{str_depth}(u)$, and note the next factor position.

After this pass, we have determined the z_W witnesses by the '1's stored in B_W . We use the witnesses in the next pass to compute the referred positions (see Figure 3).

Pass (b). We clear B_V , create a rank-support on B_W and allocate an array W consuming $z_W \lg n$ bits. We use W to map a witness id to a text position (referred positions). Having this array as a working space, $W[w]$ becomes the label of the leaf from which we visited the witness w in the first place. So we find the referred position of a referencing factor f in $W[w]$ when visiting w again from a different leaf corresponding to f . The length of f is the string depth of w . Since fresh factors consist of single characters, we can output a fresh factor by applying the *head*-function to its corresponding leaf.

Compressing B_W . Instead of directly marking nodes in B_W , we can allocate $z \lg n$ bits (we can use the space later for the array W) storing the pre-order numbers of all witnesses during Pass (a). Afterwards, we can create a compressed bit vector with constant time rank-support, representing B_W .

Algorithm 3 LZ77 Pass (b)

```
1:  $B_V.clear$ 
2:  $B_W.add\_rank\_support$ 
3:  $p \leftarrow 1$  {tracks next leaf corresponding to a factor}
4:  $z_W \leftarrow B_W.rank_1(n)$ 
5:  $W \leftarrow$  array of size  $z_W \lg n$  {maps witness ids to text positions}
6:  $\ell \leftarrow smallest\_leaf$ 
7: repeat
8:    $v \leftarrow parent(\ell)$ 
9:   while  $v \neq root$  do
10:    if  $B_V[v] = 1$  then {Invariant:  $B_V[v] = 1 \wedge p = label(\ell) \Rightarrow B_W(v) = 1$ }
11:      if  $label(\ell) = p$  then { $\ell$  corresponds to a factor}
12:        output text position  $W[B_W.rank_1(v)]$ 
13:        output factor length  $str\_depth(v)$ 
14:         $p \leftarrow p + str\_depth(v)$  {determine next starting factor}
15:      end if
16:      break
17:    end if
18:    if  $B_W[v] = 1$  then
19:       $W[B_W.rank_1(v)] \leftarrow label(\ell)$ 
20:    end if
21:     $B_V[v] \leftarrow 1$ 
22:     $v \leftarrow parent(v)$ 
23:  end while
24:  if  $label(\ell) = p$  then {We are currently processing a leaf of a fresh factor}
25:    output character  $head(\ell)$ 
26:    output factor length 1
27:    incr  $p$ 
28:  end if
29:   $\ell \leftarrow next\_leaf(\ell)$ 
30: until  $\ell = smallest\_leaf$ 
```

5.1 Trade-Off Variant

We can reduce W to $\varepsilon z \lg z$ by performing both passes $1/\varepsilon$ times. Therefore, we prematurely stop Pass (a) after counting ε many witnesses. We store the label j of the last visited leaf in order to resume Pass (a) after outputting the found factors. Since z_W is now ε , we need not modify Pass (b). When running Pass (a) again to capture the next ε witnesses (some may be the same), we suppress the marking in B_W when visiting leaves corresponding to referencing factors whose factor positions are at most j .

Since we run both passes $1/\varepsilon$ times we get $\mathcal{O}(n/\varepsilon)$ time overall.

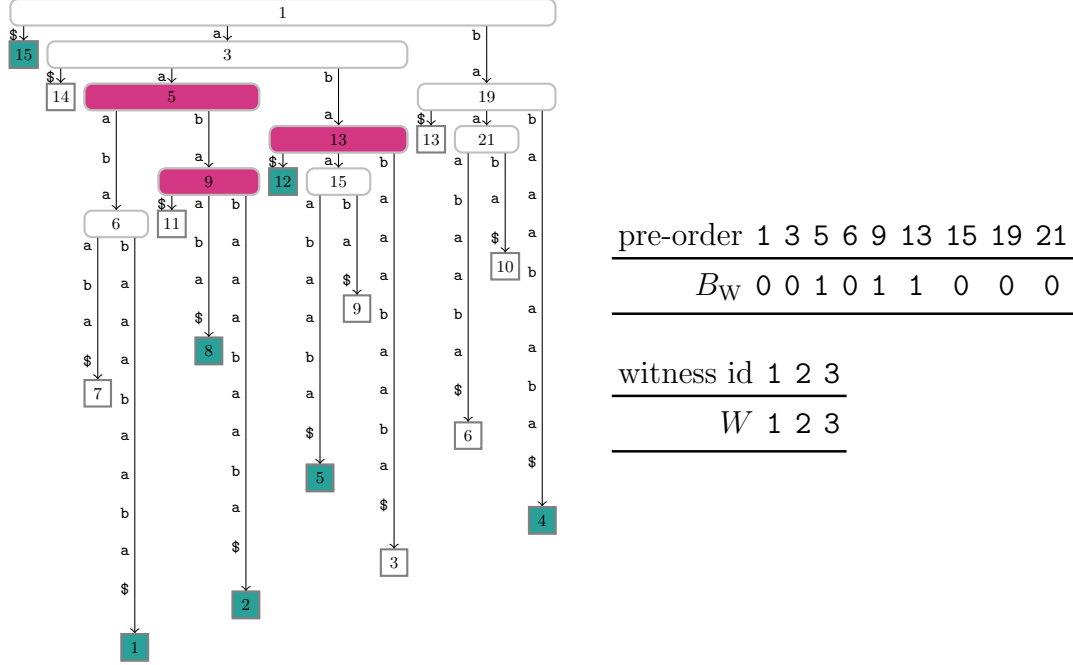


Figure 3: Our LZ77 algorithm determines the witness nodes and the leaves corresponding to factors in Pass (a). Considering our running example $T = \text{aaababaaabaaba}\$, the witness nodes are the nodes with the pre-order numbers 5, 10, and 14, and the leaves corresponding to factors have the labels 1, 2, 4, 5, 8, 12, and 15. Each witness w is the lowest ancestor of a leaf corresponding to a factor f with the property that the referred position of f is the label of a leaf contained in w 's subtree. For instance, the leaf corresponding to the 4-th factor has the label 5. Its witness has pre-order number 13, leading to the leaf with the label 3. So the referred position of the 4-th factor is 3. The length of the 4-th factor is the string depth of its witness. We show W and B_W after Pass (b). In this example, $z_W = 3$ and $z_R = 4$.$

5.2 Classic LZ77 factorization

The LZ77 and the classic-LZ77 factorization differ in the fact that a factor introduces always a new character at its end in the classic flavor. We can easily adopt our LZ77 factorization algorithm to the classic flavor. To this end, when processing a corresponding leaf, we skip additionally to the string depth of its witness one character, making the currently processed factor longer by one character. Like in the LZ78 part, we can get the new character of a referencing factor with factor index x by accessing the leaf ℓ' whose label is one text position smaller than the label of the leaf corresponding to the $(x+1)$ -th factor. We conducted the classic LZ77 factorization on our running example in Figure 4.

We additionally have to check in Algorithm 2 line 16 that we reach the root from a corresponding leaf. This is no longer an invariant: If this condition does not hold, then the previous factor consumed the new character that would be used as a fresh factor in our standard variant.




Figure 4: Applying the LZ77-classic algorithm to our running example $T = \text{aaababaaabaaba}\$$, the witness nodes have the pre-order numbers 3, 5, 13, and 15, and the leaves corresponding to factors have the labels 1, 2, 5, 9, and 14. In this example, we have $z_W = z_R = 4$.

6 LZ78

A natural representation of the LZ78 factors is a trie, the so-called **LZ trie**. Each node in the trie represents a factor and is labeled by its index. If the x -th factor refers to the y -th factor, then there is a node u having a child v such that u and v have the unique labels y and x , respectively. The edge (u, v) is labeled by the last character of the x -th factor (the newly introduced character). A node with the label x is the child of the root iff the x -th factor is a fresh factor.

The LZ trie representation is used in the algorithms presented below. While the streaming algorithm computes the topology of the LZ trie, the storing variant explicitly creates the LZ trie for querying. We can compute the factorization with $5n + z \lg z + o(n)$ additional bits of working space when streaming the output, or with $6n + z(\lg \sigma + \lg z + 3) + o(n)$ additional bits of working space when storing the LZ trie explicitly, labeling each node of the trie by a factor index.

6.1 Superimposition

The main idea is the superimposition of the suffix trie on the suffix tree, borrowed from Nakashima et al. [11]: The LZ trie is a connected subgraph of the suffix trie containing its root (see Figure 6). Regarding the suffix tree, the LZ nodes are either already represented by an ST node (explicit), or lie on an ST edge (implicit). To ease explanation, we identify each edge $e = (u, v)$ of ST uniquely with its ending node v ,

i.e., we implicitly convert between the edge e and its in-going node v (each node except **root** is associated with an edge). In order to address all LZ nodes, we keep track of how far an edge on the suffix tree got explored during the parsing. To this end, for an edge $e = (u, v)$, we define the **exploration counter** $0 \leq n_v \leq c(e)$ storing how far e is explored. If $n_v = 0$, then the factorization has not (yet) explored e , whereas $n_v = c(e)$ tells us that we have already reached v . Unfortunately, storing n_v in an integer array for all edges costs us $2n \lg n$ bits.

Our idea is to choose different representations of the exploration counters dependent on the state (not, partially or fully explored) and the number of descendants of a node. First, we mark the fully explored edges in a bit vector B_V (dynamically) such that we do not need to store their exploration counters. Further, we do not represent n_v for a node v with parent u until n_u got fully explored. Now let us focus on the rest of the nodes. We classify each node v based on the number of descendants of v , and select an explicit representation if this number is large, otherwise we maintain n_v implicitly. The classification of v is based on the following definitions borrow from [20]: If v 's subtree has at most $\lg n$ nodes, we call v **micro**. If v is not a micro node, but all its children are micro, then we call v a **jump node**. So a subtree rooted at a jump node contains at least $\lg n$ nodes, and the subtrees rooted at different jump nodes are pairwise disjoint (they do not share a node). This means that there are at most $n/\lg n$ jump nodes. We mark each jump node v in a bit vector B_J , and store n_v in an integer array J . The array J has at most $n/\lg n$ entries and therefore consumes at most $(n/\lg n) \lg n = n$ bits. Let us consider a node v that is not micro and whose in-going edge did get partially explored. If v is a jump node, then we look-up n_v in J . Otherwise, v has at least one descendant that is a jump node. Since v 's in-going edge is not fully explored, the exploration counters of all edges in the subtree rooted at v are zero. So we can abuse an exploration counter of a jump node belonging to this subtree to represent n_v until n_v gets full. For instance, we can always use the leftmost jump node of v that can be accessed by $B_J.\text{select}_1(B_J.\text{rank}_1(v) + 1)$.

The exploration counters of the micro nodes are maintained implicitly by a bit vector marking visited corresponding leaves: During a pass, when exploring a new factor on the in-going edge of a micro node v , we mark the currently accessed leaf (which will always be a leaf in the subtree rooted at v) in a bit vector B_C . By applying *popcount* to B_C , we can count how many leaves had been accessed belonging to the subtree rooted at v . This number is exactly n_v . We can compute n_v in constant time, since there are at most $\lg n$ leaves in the subtree rooted at v . After fully exploring the edge of v , we clear the area in B_C belonging to the leaves contained in v 's subtree. By doing so, the counter n_u of every micro child u of v is reset.

Dividing suffix tree nodes in explored/unexplored and not-micro/micro creates two boundaries, as illustrated in Figure 5.

Applying this procedure during a pass, we can determine the fully explored edges and collect n_v of each node v whose in-going edge got partially explored (by the definition of the jump nodes, and since we clear parts of B_C after full exploration).

6.2 Streaming Variant

We do two passes:

Algorithm 4 Function *find_edge* finds the first edge (u, v) on the path from the root to ℓ that is not yet fully explored.

Require: leaf ℓ

```

1:  $v \leftarrow \text{root}$ 
2:  $d \leftarrow 0$  {node depth}
3: repeat {find first edge on path from root to  $\ell$  that is not fully explored}
4:   incr  $d$ 
5:    $v \leftarrow \text{level\_anc}(\ell, d)$ 
6: until  $v = \ell$  or  $B_V[v] = 0$ 
7:  $u \leftarrow \text{parent}(v)$  {new factor is on the edge  $(u, v)$ }
8: return  $(u, v)$ 

```




Figure 5: Our LZ78 algorithm divides ST by two boundaries: Nodes having at least $\lg n$ nodes in their subtree (hatched upper cone), and (partially) explored edges belonging to the LZ trie (colored upper cone). The nodes whose exploration counters are stored implicitly are directly below the fringe of the LZ trie. During the passes, we always search for an ST edge (u, v) crossing the boundary of the already discovered part.

- (a) create B_W so we can address the witnesses (see Algorithm 6), and
- (b) stream the output by using a helper array mapping witness ids to factor indices (see Algorithm 7).

We explain the passes in detail, after introducing their commonality and a helpful lemma:

LZ78 Passes. Since referencing factors address factor indices (z options) instead of text positions (n options), we are only interested in the leaves corresponding to a factor. Starting with *smallest_leaf*, which corresponds to the first factor, we can compute the length of the factor corresponding to the currently accessed leaf so that we know the distance (in text positions) to the next corresponding leaf.

Lemma 6.1 ([12, Lemma 4]). *Let $e = (u, v)$ be an ST edge, and u the parent of the node v . Then $n_v \leq \min(c(e), s)$, where s is the number of leaves of the subtree rooted at v .*

Pass (a). The main goal of this pass is to determine the topology of the LZ trie with respect to the superimposition. Starting with an LZ trie consisting only of the root, we build the LZ trie successively by filling up the exploration counters. If the exploration counter of an edge is filled up, we mark its in-going node in the bit vector B_V .

Algorithm 5 Function *explore* post-increments n_v , returning its previous value.

Ensure: $B_J[v] \leftarrow 1$ for every jump node v

Ensure: $J \leftarrow$ integer array with $n/\lg n$ entries, i.e., n bits.

Require: node v , $s = \text{str_depth}(v)$

```

1: if  $v$  is a jump node then
2:    $m \leftarrow J[B_J.\text{rank}_1(v)]$ 
3: else if  $v$  is not micro then
4:    $m \leftarrow J[B_J.\text{select}_1(B_J.\text{rank}_1(v) + 1)]$ 
5: else
6:    $m \leftarrow \text{popcount}(B_C[l\text{most\_leaf}(v), r\text{most\_leaf}(v)])$ 
7: end if
8: Let  $\ell, \ell'$  be two leaf whose LCA is  $v$ 
9:  $\ell = (\text{next\_leaf})^{m+s+1}[\ell]$ 
10:  $\ell' = (\text{next\_leaf})^{m+s+1}[\ell']$ 
11: if  $\text{head}(\ell) \neq \text{head}(\ell')$  then {edge  $(u, v)$  now fully explored}
12:    $B_V[v] \leftarrow 1$  {set  $v$  as fully explored}
13:    $B_C.\text{clear}[l\text{most\_leaf}(v), r\text{most\_leaf}(v)]$  {reset the counting so that we can work
    with the children of  $v$ }
14: else if  $v$  is micro then {edge  $(u, v)$  has at least one character unexplored}
15:    $B_C[\ell] \leftarrow 1$  {increment  $n_e$  for the micro node  $v$ }
16: end if
17: if  $B_J[v] = 1$  then {increment  $n_e$  for the jump node  $v$ }
18:   incr  $J[B_J.\text{rank}_1(v)]$ 
19: else if  $v$  is not micro then
20:   incr  $J[B_J.\text{select}_1(B_J.\text{rank}_1(v) + 1)]$ 
21: end if
22: return  $m$ 

```

Assume that we visit a leaf ℓ . We want to find the first edge on the path from **root** to ℓ that is either unexplored or partially explored. By invoking level ancestor queries, we traverse from the root to an edge $e = (u, v)$, where $n_v < c(e)$ and u is (already) represented as a node in the LZ trie. If v is an internal node with $n_v = 0$, we make v a witness by marking v in B_W (the idea is that the edge e is *superimposed* by some LZ nodes).

Regardless that, we add a new factor by incrementing n_v . If the edge e now got fully explored, we additionally mark v in B_V . Whether the edge e got fully explored, can be determined with the *next_leaf* function: First, if v is a leaf, the edge (u, v) can be explored at most once (by Lemma 6.1). Otherwise, we choose a leaf ℓ' such that the LCA of ℓ and ℓ' is v . The idea is that $\text{str_depth}(v)$ is the length of the longest common prefix of two suffixes corresponding to two leaves (e.g., ℓ and ℓ') having v as their LCA. So we can compare the m -th character of both respective suffixes by applying *next_leaf* m -times on both leaves before using the *head*-function. With $m := \text{str_depth}(u) + n_v + 1$ we can check whether the edge (u, v) got fully explored. Additionally, we can determine the label of the next corresponding factor. Although

Algorithm 6 LZ78 Pass (a)

```
1:  $\ell \leftarrow \text{smallest\_leaf}$ 
2: repeat
3:    $(u, v) \leftarrow \text{find\_edge}(\ell)$ 
4:    $s \leftarrow \text{str\_depth}(u)$ 
5:   if  $v = \ell$  then
6:      $\ell = (\text{next\_leaf})^{s+1}[\ell]$ 
7:     continue
8:   end if
9:    $B_W[v] \leftarrow 1$   $\{v \text{ is a witness}\}$ 
10:   $m \leftarrow \text{explore}(v, \ell)$   $\{m \leftarrow n_v; n_v \leftarrow n_v + 1\}$ 
11:   $\ell = (\text{next\_leaf})^{m+s+1}(\ell)$ 
12: until  $\ell = \text{smallest\_leaf}$ 
```

we apply *next_leaf* as many times as the factor length, we still get linear time overall, because concatenating all factors yields the text T .

Pass (b). This pass is nearly identical to Pass (a). We explore the LZ trie nodes again, but this time we already have the witnesses. So we keep B_W , but reset the exploration counters and B_V .

For finding the referred indices, we create an array W with $z_W \lg z$ bits to store a factor index for each witness id. The witness ids are determined by B_W . The factor indices are given by a counter variable tracking the number of visited corresponding leaves, i.e., the number of processed factors.

Assume that we visit the leaf ℓ corresponding to the x -th factor, i.e., ℓ is the x -th visited corresponding leaf. Again by level ancestor queries, we determine the edge $e = (u, v)$ on the path from the root to ℓ , where u is in B_V and v not.

If v is an internal node, then v is a witness. In this case, we retrieve $y := W[B_W.\text{rank}_1(v)]$. If y is defined, then the x -th factor refers to the y -th factor.

If y is undefined (i.e., its value has not yet been initialized), or if v is a leaf (i.e., $v = \ell$), then the x -th factor is either a fresh factor if v is a child of **root**, or the x -th factor refers to $W[B_W.\text{rank}_1(\text{parent}(v))]$.

If v is an internal node, we set $W[B_W.\text{rank}_1(v)] \leftarrow x$, and increment n_v (thus exploring the LZ trie like before). Like before, when v 's in-going edge gets fully explored, we mark v in B_V .

So far, we can output the referred index of the x -th factor, if it exists. We get the new character of the x -th factor (i.e., the last letter of the factor) by accessing the leaf ℓ' that is (with respect to text order) *before* the leaf corresponding to the $(x + 1)$ -th factor; then we can output the new character by $\text{head}(\ell')$.

6.3 Explicitly storing the LZ trie

Our goal is to represent the LZ78 factorization by three data structures: The first one is a balanced parenthesis sequence storing the LZ trie topology. It is accompanied by an array storing the factor index and an array storing the ending character, for each LZ node. We can build the balanced parenthesis sequence directly after Pass (a): To

Algorithm 7 LZ78 Pass (b)

```
1:  $B_W.add\_rank\_support$ 
2:  $z_W \leftarrow B_W.rank_1(n)$ 
3:  $W \leftarrow$  array of size  $z_W \lg z$  bits
4:  $x \leftarrow 0$ 
5:  $\ell \leftarrow smallest\_leaf$ 
6: repeat
7:    $(u, v) \leftarrow find\_edge(\ell)$ 
8:    $s \leftarrow str\_depth(u)$ 
9:   if  $B_W[v] = 0$  then
10:    if  $v$  is a child of root then
11:      the  $x$ -th factor is a fresh factor
12:    else
13:      the  $x$ -th factor refers to  $W[B_W.rank_1(parent(\ell))]$ 
14:    end if
15:  else
16:    the  $x$ -th factor refers to  $W[B_W.rank_1(v)]$ 
17:  end if
18:  incr  $x$ 
19:  if  $v = \ell$  then
20:     $\ell = (next\_leaf)^{s+1}[\ell]$ 
21:    continue
22:  end if
23:   $m \leftarrow explore(v, \ell) \{m \leftarrow n_v; n_v \leftarrow n_v + 1\}$ 
24:   $\ell = (next\_leaf)^{m+s}(\ell)$ 
25:  output character  $head(\ell)$  belonging to the  $(x - 1)$ -th factor
26:   $\ell = next\_leaf(\ell)$ 
27: until  $\ell = smallest\_leaf$ 
```

this end, we traverse the suffix tree with a depth first search starting at the root. While traversing **ST** we create $c(e)$ LZ nodes for each edge $e = (u, v)$ if v belongs to B_V , or create n_v LZ nodes otherwise. We can retrieve the value of n_v or $c(e)$ because of the following: For the former, we know the exploration counters of all *partial* discovered edges during Pass (a). For the latter, we can recompute the $c(\cdot)$ values of the *fully* discovered edges with str_depth without worsening the linear time bound.

During the traversal of **ST** we can build the $z \lg \sigma$ -bits array storing the ending characters: Assume that we create the LZ nodes u' and v' on the **ST** edge (u, v) , where u' is the parent of v' . We can obtain the label of the edge (u', v') consisting of a character by invoking $next_leaf$ and $head$: Given that s is the string depth of v' in the LZ trie, applying $next_leaf$ s times to a leaf ℓ in the subtree rooted at v returns a leaf ℓ' whose $head(\ell')$ -value is the label in question.

Overall, we can compute the LZ trie in $\mathcal{O}(n)$ time, storing its BP representation along with the edge labels in $2z + z \lg \sigma + o(z)$ bits.

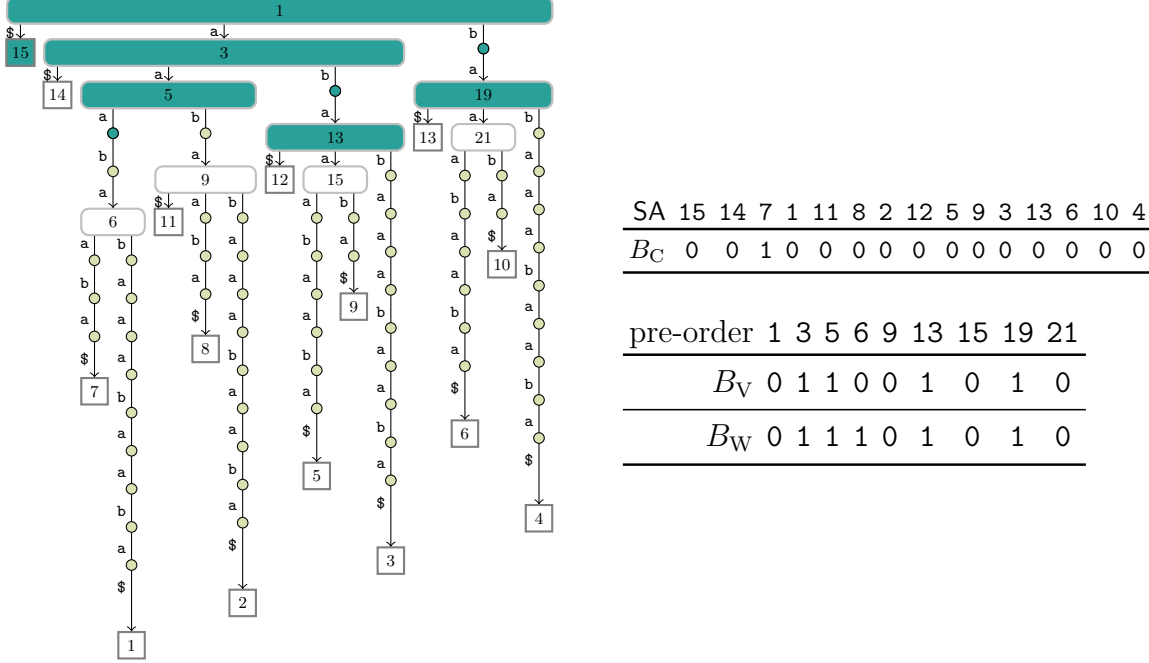


Figure 6: We can get the suffix trie (conceptually) by exchanging every ST edge e with $c(e) - 1$ new suffix trie nodes superimposing e . These new suffix trie nodes are the small rounded nodes in the depicted tree. They represent the implicit suffix trie nodes, while the remaining ST nodes represent the explicit suffix trie nodes. Dark colored and hatched nodes represent the nodes of the LZ78 trie. We show B_C , B_V , and B_W after Pass (a).

Annotating the LZ nodes with the factor indices cannot be done in linear time without a pre-computation step, which is topic of this paragraph. Here, we want to compute a mapping between suffix tree nodes and LZ nodes such that we can access the LZ node corresponding to the currently explored factor during a pass in constant time. To this end we create a bit vector B_U marking suffix tree nodes and a bit vector B_E marking LZ nodes. The former marks all suffix tree nodes v with $n_v > 0$, the latter marks all LZ trie children of every explicit LZ node. The last thing we do is adding a rank-support to B_U , and a select-support to B_E . Now assume that there is a suffix tree node v marked in B_U . The LZ trie contains a node v' marked with B_E such that $v' = B_E.\text{select}_1(B_U.\text{rank}_1(v))$ (conceptually $B_E.\text{rank}_1(v') = B_U.\text{rank}_1(v)$). The node v' has the following properties: First, the LZ trie parent of v' is an explicit LZ node that is represented by the suffix tree parent of v (v has a parent since the root is not marked in B_U). Second, there are $n_v - 1$ nodes below v' forming a unary tree (also called linear graph or path graph). Due to the balanced parenthesis representation of the LZ trie, the query $v'.\text{descendant}(j)$ returning the descendant of v' at depth $\text{depth}(v') + j$ for $0 \leq j \leq n_v - 1$ can be answered in constant time.

Now we explain the final pass. We exchange Pass (b) with

(c) store the factor index and the ending character of each LZ node in an array W .

Pass (c). We create an array W with $z \lg z$ bits to store a factor index. We recompute B_V and the exploration counters. We count the current factor index with a

variable x . Assume that we visit the leaf ℓ corresponding to the x -th factor. By level ancestor queries, we find the edge $e = (u, v)$ in **ST** on the path from **root** to ℓ , where u belongs to B_V , but v not. We set $W[B_E.\text{select}_1(B_U.\text{rank}_1(u)).\text{descendant}(n_v)] \leftarrow x$.

Acknowledgements

We thank Veli Mäkinen for outlining us the differences between [6] and [8], and Johannes Fischer for some helpful comments. Most parts of this work was done during a visit at the graduate school of information science and technology of the university of Tokyo, supported by the *Studienwerk für Deutsch-Japanischen Kulturaustausch in NRW e.V.*

References

- [1] J. Ziv and A. Lempel, “A Universal Algorithm for Sequential Data Compression,” *IEEE Transactions on Information Theory*, vol. 23, no. 3, pp. 337–343, 1977.
- [2] —, “Compression of Individual Sequences via Variable-Rate Coding,” *IEEE Transactions on Information Theory*, vol. 24, no. 5, pp. 530–536, 1978.
- [3] D. Kosolobov, “Faster Lightweight Lempel-Ziv Parsing,” in *MFCS*. Springer Berlin Heidelberg, 2015, vol. 9235, pp. 432–444.
- [4] J. Kärkkäinen, D. Kempa, and S. J. Puglisi, “Lightweight Lempel-Ziv Parsing,” in *Experimental Algorithms*, 2013, pp. 139–150.
- [5] D. Belazzougui and S. J. Puglisi, “Range Predecessor and Lempel-Ziv Parsing,” *SODA*, pp. 2053–2071, 2016.
- [6] D. Belazzougui, “Linear Time Construction of Compressed Text Indices in Compact Space,” in *STOC*. ACM, 2014, pp. 148–193.
- [7] M. Burrows and D. J. Wheeler, “A Block-sorting Lossless Data Compression Algorithm,” Digital Equipment Corporation, Tech. Rep., 1994.
- [8] D. Belazzougui, “Linear time construction of compressed text indices in compact space,” *ArXiv CoRR*, vol. abs/1401.0936, 2015.
- [9] J. Fischer and P. Gawrychowski, “Alphabet-Dependent String Searching with Wexponential Search Trees,” in *CPM*, 2015, pp. 160–171.
- [10] J. Jansson, K. Sadakane, and W. Sung, “Linked Dynamic Tries with Applications to LZ-Compression in Sublinear Time and Space,” *Algorithmica*, vol. 71, no. 4, pp. 969–988, 2015.
- [11] Y. Nakashima, T. I, S. Inenaga, H. Bannai, and M. Takeda, “Constructing LZ78 Tries and Position Heaps in Linear Time for Large Alphabets,” *Inform. Process. Lett.*, vol. 115, no. 9, pp. 655 – 659, 2015.
- [12] J. Fischer, T. I, and D. Köppl, “Lempel-Ziv Computation in Small Space (LZ-CISS),” in *CPM*, 2015, pp. 172–184.
- [13] J. I. Munro, “Tables,” in *Proc. FSTTCS*, ser. LNCS, vol. 1180. Springer, 1996, pp. 37–42.
- [14] R. Raman, V. Raman, and S. R. Satti, “Succinct Indexable Dictionaries with Applications to Encoding K-ary Trees, Prefix Sums and Multisets,” *ACM Trans. Algorithms*, vol. 3, no. 4, 2007.

- [15] G. Franceschini, S. Muthukrishnan, and M. Pătraşcu, “Radix Sorting with No Extra Space,” in *ESA*. Springer, 2007, vol. 4698, pp. 194–205.
- [16] R. Grossi and J. S. Vitter, “Compressed Suffix Arrays and Suffix Trees with Applications to Text Indexing and String Matching,” *SIAM Journal on Computing*, vol. 35, no. 2, pp. 378–407, 2005.
- [17] G. J. Jacobson, “Space-efficient static trees and graphs,” in *Proc. FOCS*. IEEE Computer Society, 1989, pp. 549–554.
- [18] K. Sadakane, “Compressed Suffix Trees with Full Functionality,” *Theory of Computing Systems*, vol. 41, no. 4, pp. 589–607, 2007.
- [19] G. Navarro and K. Sadakane, “Fully Functional Static and Dynamic Succinct Trees,” *ACM Trans. Algorithms*, vol. 10, no. 3, pp. 16:1–16:39, 2014.
- [20] K. Sadakane and R. Grossi, “Squeezing succinct data structures into entropy bounds,” in *Proc. SODA*. ACM/SIAM, 2006, pp. 1230–1239.

A Appendix

A.1 Overview of used data structures

While describing both factorization algorithms, we used several data structures, among others bit vectors, some with rank or select-support, to achieve the small space bounds. We give here an overview (see also Table 1). We denote bit vectors with B_α for some letter α .

We use for both factorizations

- B_W marking all witness nodes,
- the array W mapping witness ids to
 - (LZ77) text positions, or
 - (LZ78) factor indices.

In LZ77 we use

- B_V marking visited nodes, and

In LZ78 we use

- B_C marking corresponding leaves (is used as a counter),
- B_J marking jump nodes
- B_V marking suffix tree nodes represented in the LZ trie (their ingoing edges got fully explored),
- the array J storing $n/\lg n$ numbers binary,
- B_E marking explicit LZ trie nodes, and
- B_U as a copy of B_V .

Common						
Name	bits				rank	select compress
B_W	$H_0(n, z)n + o(n)$				○	○
ST	$4n + o(n)$					
ψ	$\mathcal{O}(n \lg \sigma)$					
LZ77						
Name	bits	(a)	(b)		rank	select compress
B_V		○	○			
W	$z \lg n$	○	○			
LZ78						
Name	bits	(a)	(b)	(c)	rank	select
B_C		○	○	○		
B_J	$n + o(n)$	○	○	○	○	○
B_V		○	○	○		
B_E	$z + o(z)$			○		○
B_U				○		
J		○	○	○		
W	$z \lg z$	○	○	○		

Table 1: Additional data structures used while computing a LZ77/78 factorization. The letters written in brackets represent a pass (e.g., (a) refers to Pass (a)). The number of bits is omitted if it is exactly n . Circles symbolize that the data structure is used during a pass, or that it is used with a rank or select structure, or that the data structure is compressible.